

Advances in Intelligent Platform Management

**Tom Slaight
Server Management Architect
Enterprise Platforms Group
Intel Corporation**



Agenda

- **IPMI Update**
- **New system directions for IPMI**
- **New components for IPMI**
- **IPMI future directions**
- **Summary**



Itanium® is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States or other countries."

Intel
Developer
Forum
Spring 2003

Introduction

- **Audience:**
Architects, Technical Managers, Firmware Leads, and Hardware Designers
 - Involved in architecture, component selection, debug, test, or design of server baseboard and peripheral management subsystems
- **Focus: IPMI-based implementations**
 - Hardware and software components

IPMI

Intelligent Platform Management Interface

- Defines a standardized, abstracted, message-based interface to intelligent platform management hardware
- Defines standardized records for describing platform management devices and their characteristics

Promoters:



Adopters: 145 and growing

<http://developer.intel.com/design/servers/ipmi>



Initiative News

United States Home | Select a Location | Contact Us | About Intel

Products | Support

Home Computing | Business | Developer | Reseller

Advanced Search

Developer

- Hardware Design
- Software Development
- Download Software/Drivers
- Design Support/Services
- Events, Training & Pubs
- R&D/Initiatives

server home

Server Building Blocks

Itanium™ Processor Family

Hardware Developer's Resource Center

Software Developer's Resource Center

Industry Technology Specifications

Community

Related Sites

Tools & Resources

Intelligent Platform Management Interface

New IPMI v1.5 Conformance Test Suite (ICTS) Prototype 5.02 (updated on 8/05/02):

Reference Driver for IA-64 and IA-32 under Windows* .NET/2000 OS

Reference driver implementation for IA-64 and IA-32 under Windows* .NET/2000 OS available to IPMI adopters only.

- IPMI v1.5 Conformance Test Suite (ICTS) Prototype 5.02**
Includes IPMI v1.0 and IPMI v1.5 automated conformance tests, IPMI v1.5 CMDTOOL for manual IPMI v1.5 testing, support for PCI* card based IPMB and SMBus testing, and support for IPMI v1.5 new interfaces including LAN, Serial and SMBus. ICTS 5.02 is an update to ICTS 5.01 and adds new tests for IPMI 1.5 commands and includes some bug fixes as well.

developer.intel.com/design/servers/ipmi



* Other names and brands may be claimed as the property of others.




Agenda

- IPMI Update
- **New system directions for IPMI**
- New components for IPMI
- IPMI future directions
- Summary



New System Directions for IPMI

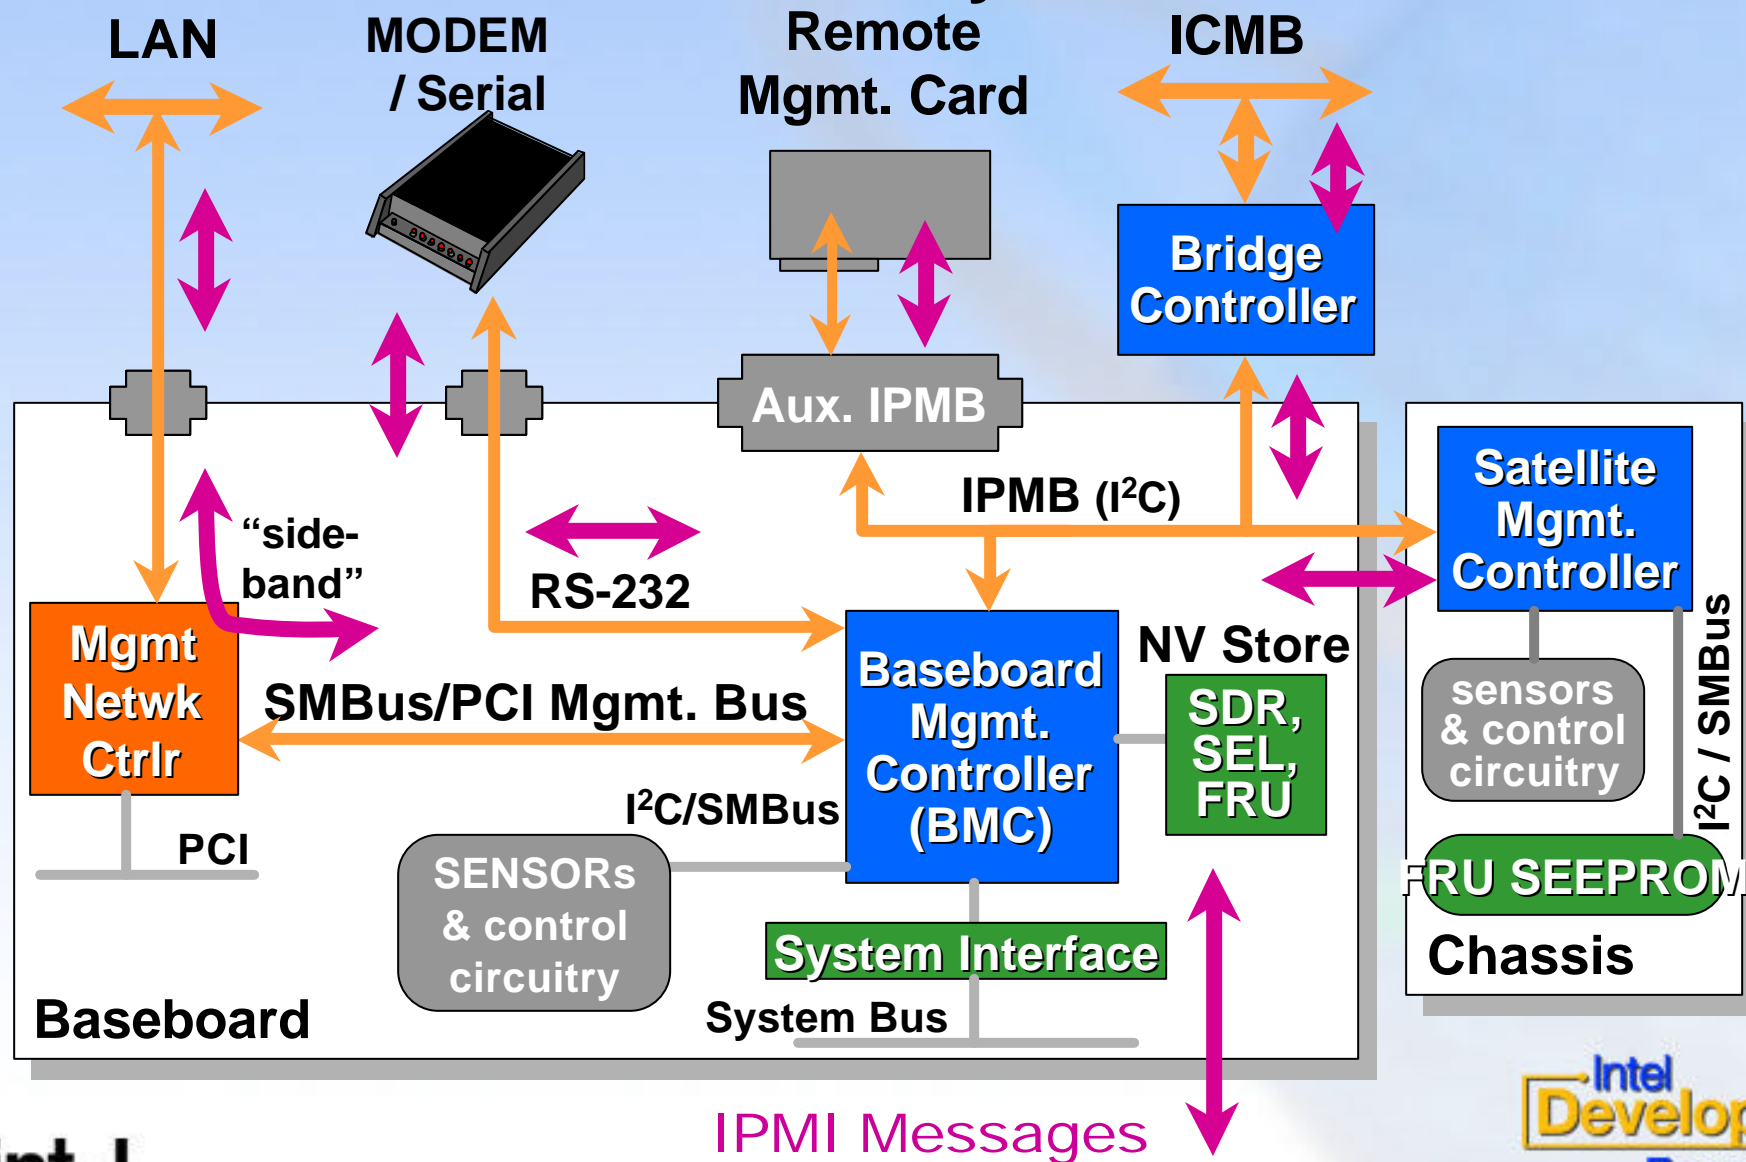
- **Modular Systems**
 - Advanced TCA & General Modular
- **Low Cost Systems**
 - Low cost BMCs
 - Add-in BMCs
- **Group Managed Systems**
 - IPMI for Group Chassis Control



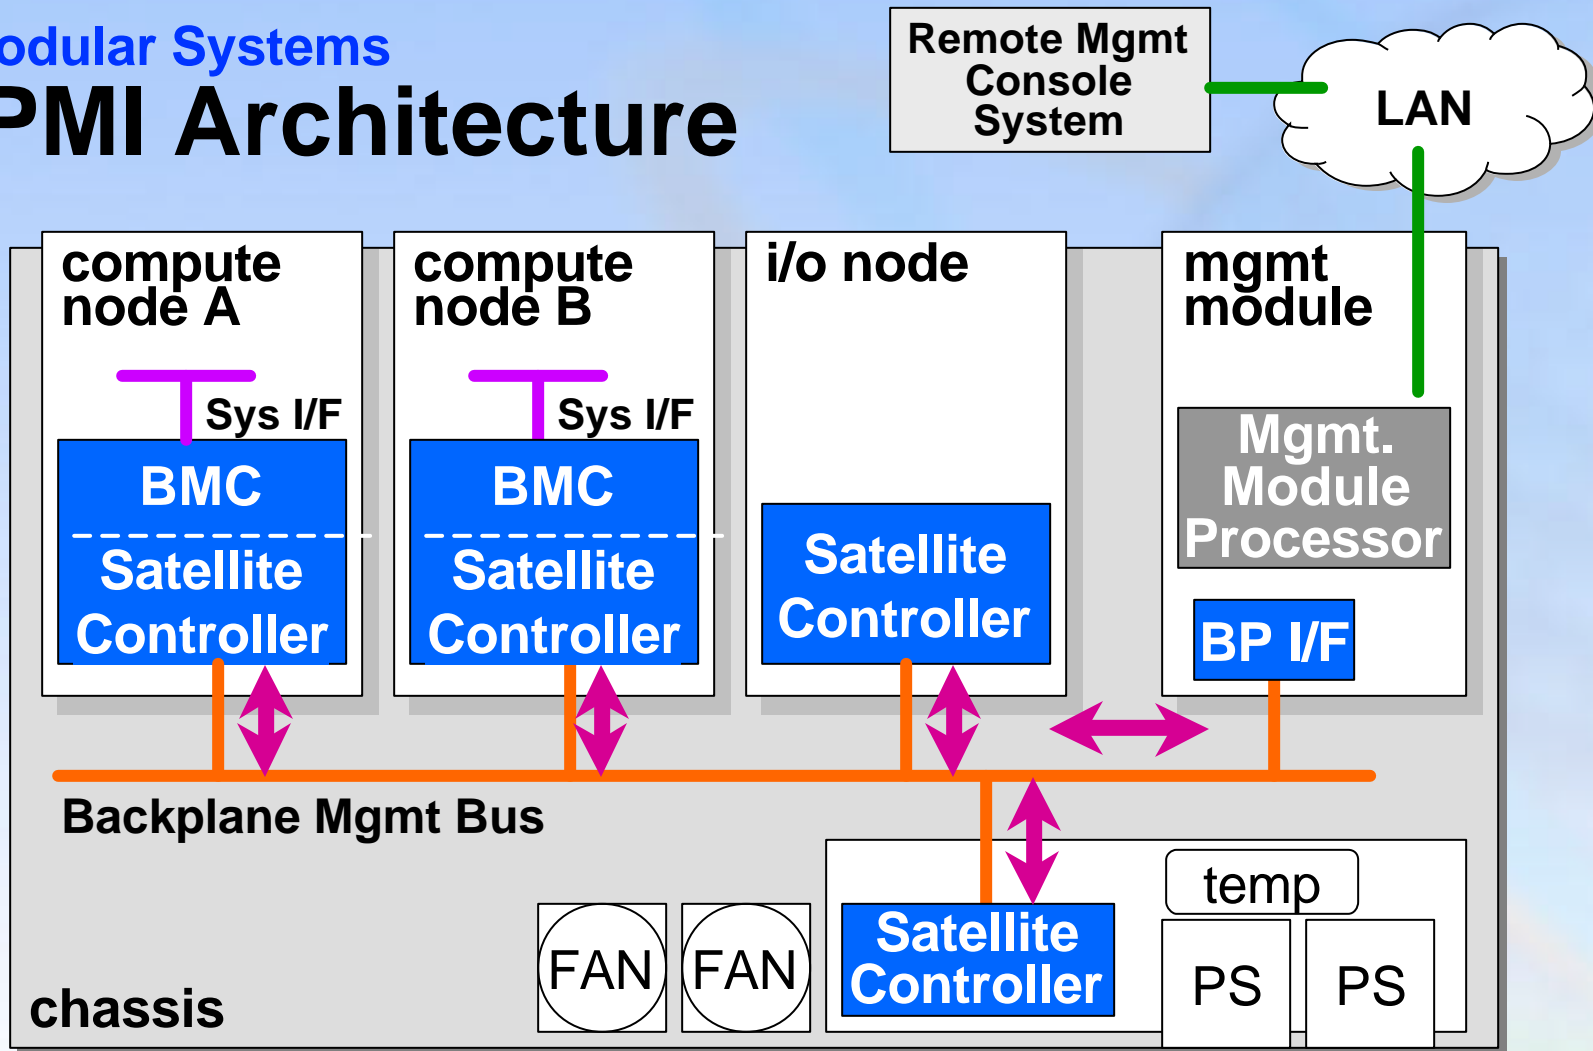
IPMI grows to enable competitive features across server classes

IPMI in modular architecture

IPMI v1.5 for Standalone Systems



Modular Systems IPMI Architecture



IPMI Messages ↔




More info in MODS135: "Using IPMI Platform Management In Modular Computer Systems"



Modular Systems

IPMI Spec Extensions

- Entity information extensions to identify node location in a given chassis
- Options to coordinate Node/FRU removal and replacement
- Management Bus 'failover' status
 - Supports redundant management bus connections



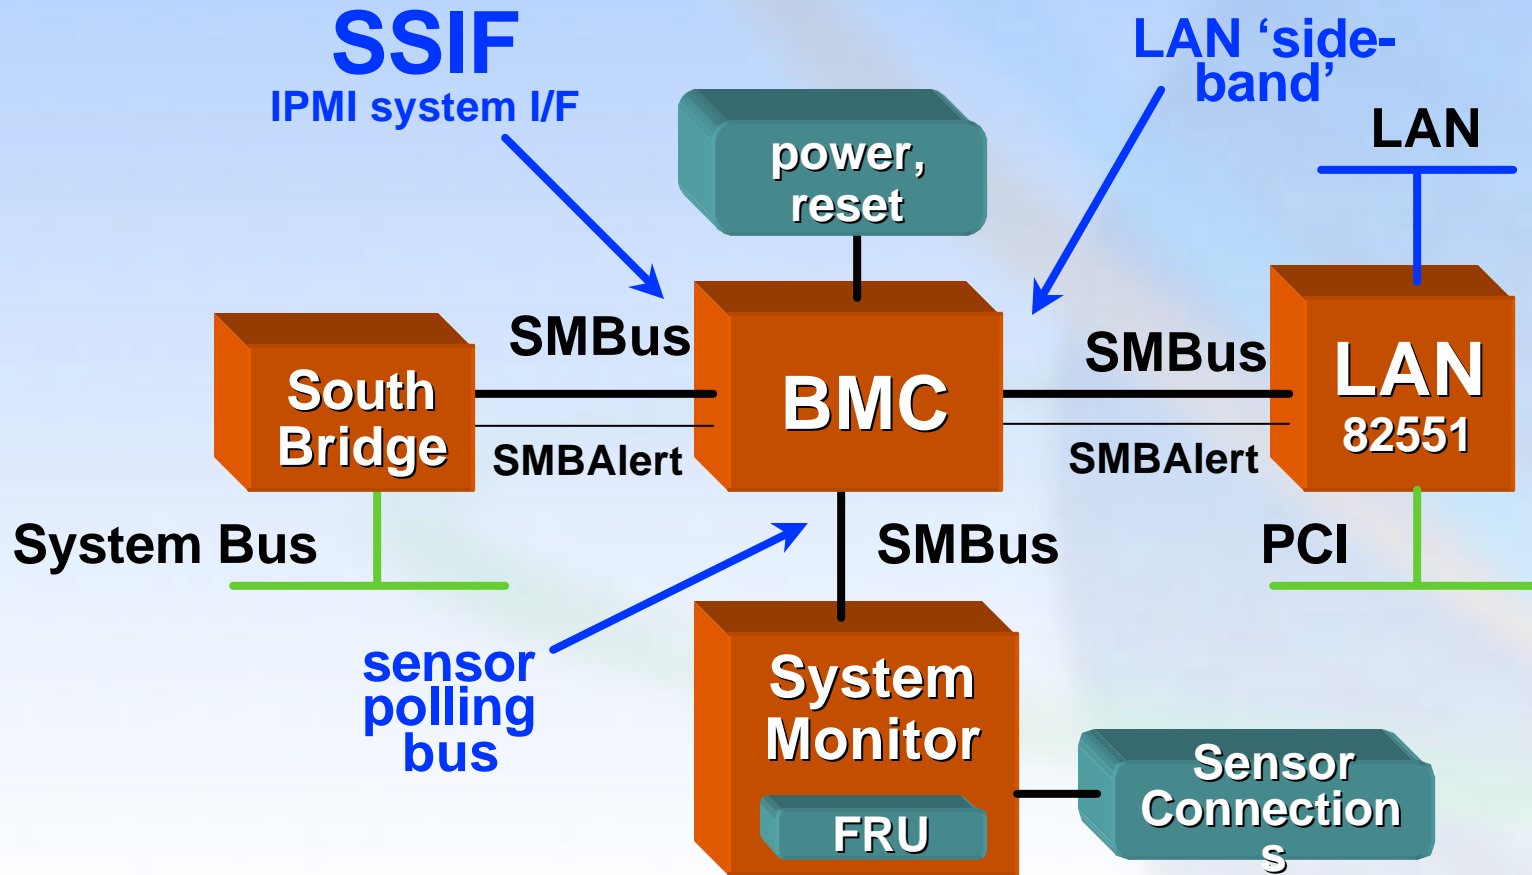
Low-cost systems

Baseline 'mini' BMCs

- **IPMI Conformant 'Baseline' BMC**
 - Covers mandatory IPMI BMC features
- **'Turnkey' firmware with limited customization and chassis support**
 - "NVRAM" configurable to motherboard sensors
 - Minimized support for IPMI options
 - E.g. May be 'LAN only'
 - May have limited options such as # of Users
- **Small, low pin-count packaging...**

Low-cost Systems

SSIF - SMBus System Interface



Low-cost Systems

SSIF - SMBus System Interface

- **Provides BMC system interface access via SMBus**
 - Low pin count
 - Relies on controller-specific SMBus drivers
- **BMC accessed as SMBus Slave device**
 - Requests delivered using 'Block Write' protocol
 - Responses retrieved using 'Block Read' protocol
 - SMBAlert signal status change/message available
 - 'Get Status' command allows interface status to be polled
 - Includes ability to write/read more than 32 data bytes using two SMBus transactions
- **New 'Reserve Device' command**
 - Optional command directs BMC to suspend access to specified device for xx milliseconds
 - Resolves access to 'shared' SMBus devices

Low-cost Systems

SSIF - Message Formats

Write_Message:

Sends encapsulated IPMI message data to controller

Format: Slave Addr | 0b, CMD=Write_Msg, Length,
Message Data, **CHECKSUM**

Read_Message:

Retrieves encapsulated IPMI message data from controller

Format: Slave Addr | 0b, CMD=Read_Msg, Slave Addr | 1b,
Length, Message Data, **CHECKSUM**

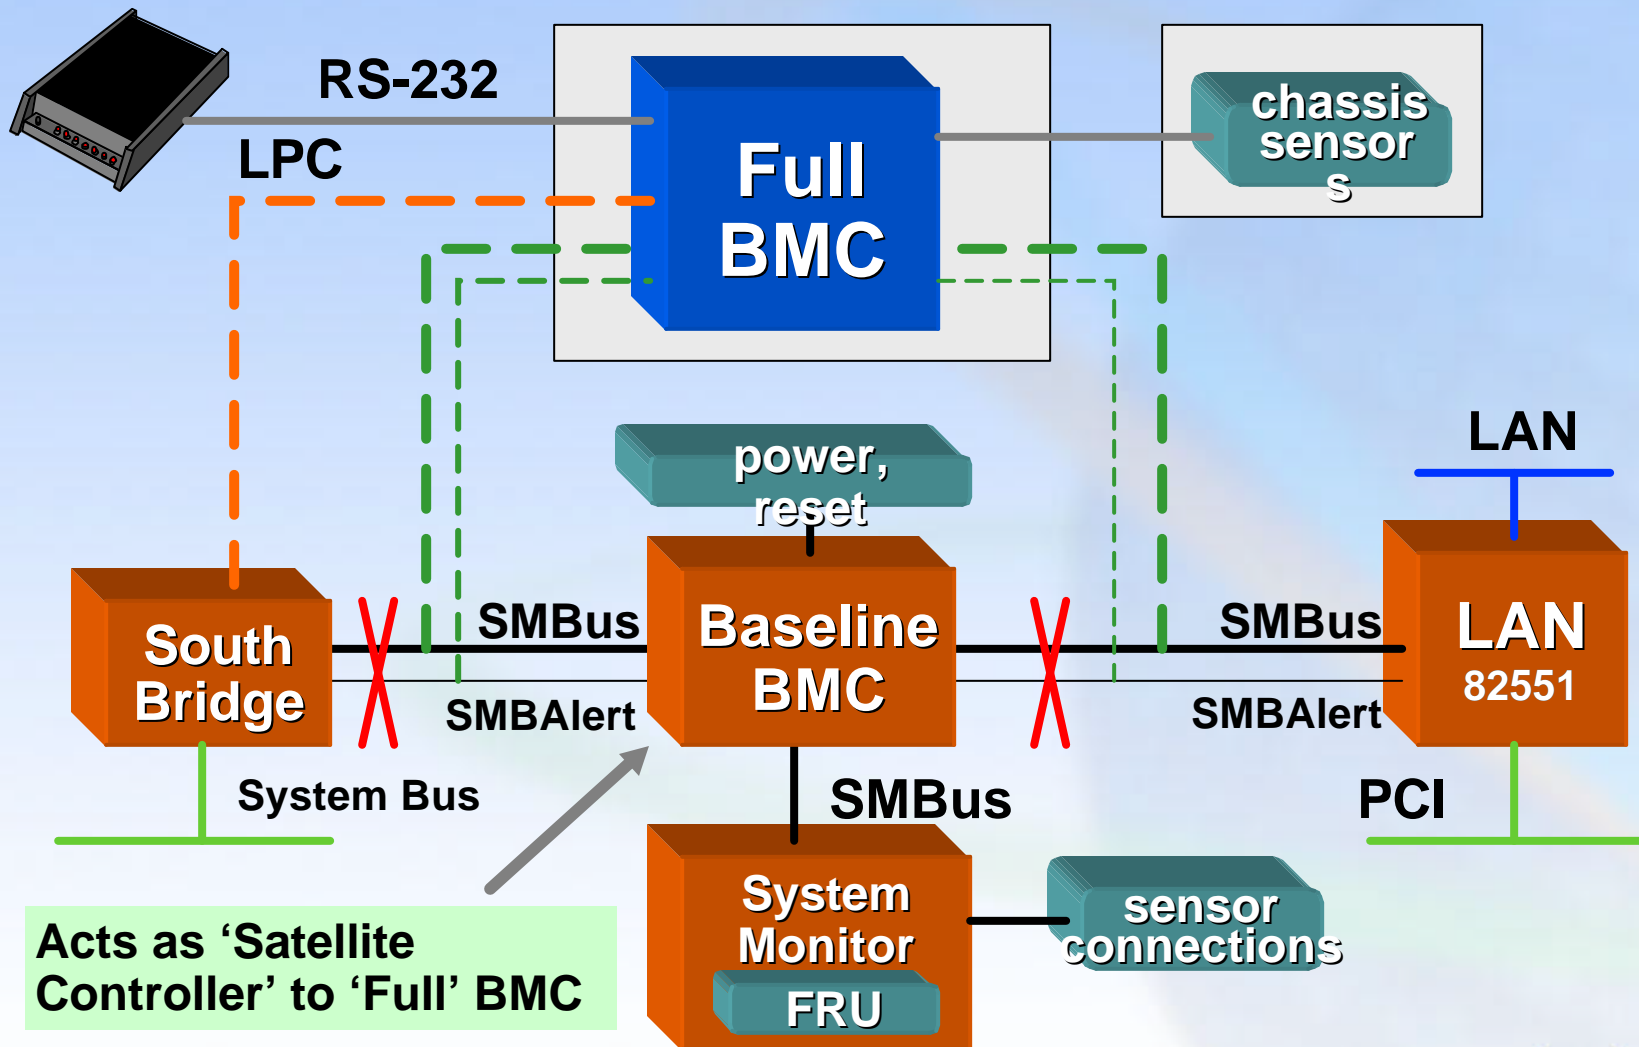
Get_Status:

Indicates message data available for reading, or error/busy status
of interfaces

Format: Slave Addr | 0b, CMD=Get_Status, Slave Addr | 1b,
Length, Message Data, **CHECKSUM**

Low-cost systems

Multi-level BMC Options




Low Cost Options enable IPMI for all Server classes



Group Managed Systems

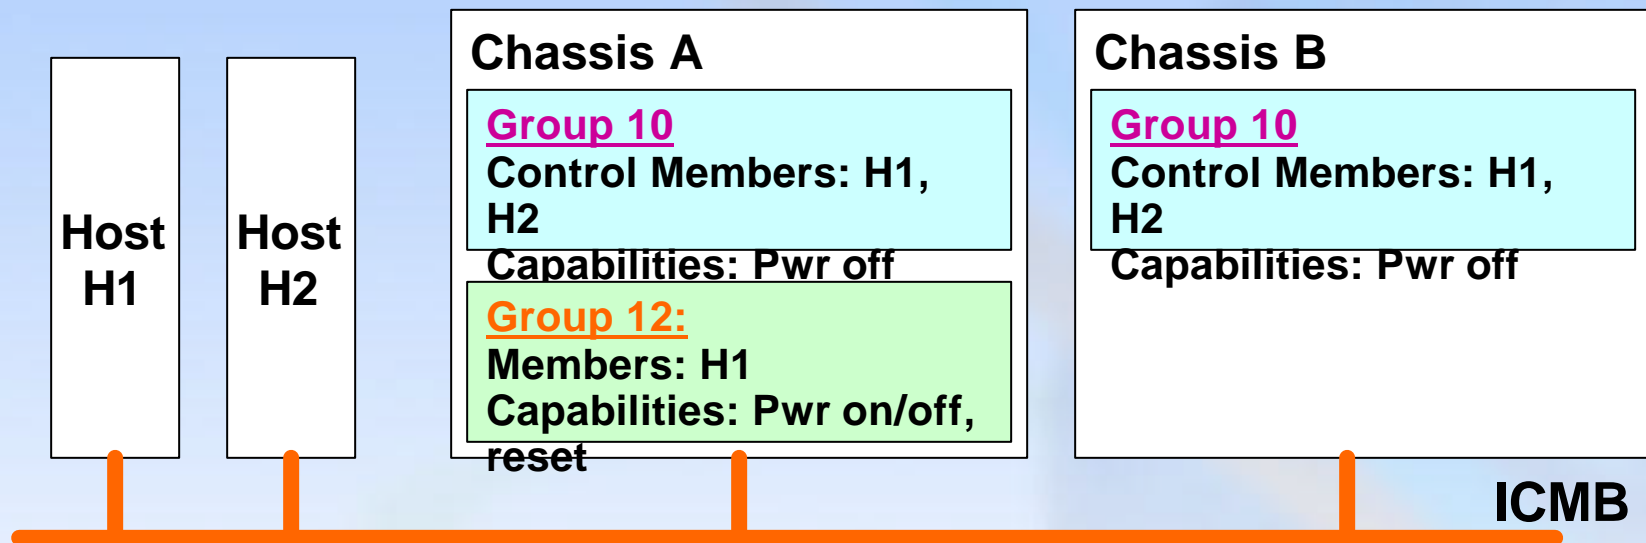
ICMB Group Chassis Control

- **Chassis Control single broadcast**
 - power on/off, reset, diagnostic interrupt
- **Group-specific operation enables, e.g.**
 - “Group 1” enabled for power-on, -off, reset
 - “Group 2” enabled for power-on only
- **Enables Power/reset sequencing**
- **Enables ‘one button’ power on**
 - e.g. pressing power button on compute chassis automatically powers up associated peripheral chassis
- **Enables ‘service lockouts’**
 - e.g. powering down a chassis blocked until all ‘controlling members’ have requested same power state



Group Management Systems

ICMB Group Chassis Control



- Up to 255 different control groups per ICMB
- Each chassis can belong to up to 4 control groups
- A single 'Group Chassis Control' command can be targeted to 4 different control groups
- Each group has up to 16 'controlling members'
- Members can 'request' or 'force' the control state

Agenda


- IPMI Update
- New system directions for IPMI
- **New components for IPMI**
- IPMI future directions
- Summary



Components

National Semiconductor “Mini” BMC

- **PC87431M*** - Targets IPMI LAN remote mgmt.
 - supports monitoring by local mgmt. s/w via SMBus
 - configurable sensor polling
- **Internal FLASH, RAM, NVRAM**
 - up to 512 bytes NV available for OEM use
- **Authenticated IPMI LAN support for:**
 - System reset, SMI/NMI, and power control
 - settable ‘Boot Options’
 - FRU, System Event Log, and SDR access
 - Sensor access
 - Alerting via IPMI/PET SNMP Traps
- **Platform Event Filtering**
 - configurable actions on events
 - power control, reset, fault light, NMI/SMI, and alert



* Other names and brands may be claimed as the property of others.

Components

General Purpose / Blade BMCs

- **Hitachi H8S/2168* ‘Single-chip’ BMC**
 - 256KB in-system programmable FLASH and 40KB SRAM on-chip
 - Six master-slave 400KHz I2C ports
 - Supports up to 3 KCS LPC channels with BT or SMIC
 - 8 A-D, 4 PWM and 16 fan tach inputs, 3 serial ports, up to 115 GP I/O
 - Supports in-system JTAG debug plus optional full capability ICE
 - 16x16mm TQFP-144 package
 - H8S/2168 IPMI Demo board available
 - Production: Now
- **Hitachi H8S/2145* ‘Single-chip’ BMC / Satellite Controller**
 - 128KB in-system programmable FLASH and 8KB SRAM on-chip
 - Two master-slave 400KHz I2C ports
 - Supports 3 KCS LPC channels
 - 8 A-D, 2 PWM and 4 fan tach inputs, 3 serial ports, up to 75 GP I/O
 - Full capability ICE debug is available
 - 14x14mm TQFP-100 package
 - Sampling: Now, Production: May



* Other names and brands may be claimed as the property of others.

Components

IPMI Software

IPMI Web Site

- **Reference drivers**
 - for Itanium™ Architecture and IA-32 under Windows* .NET/2000 and Linux
- **IPMI Conformance Test Suite**
 - Serves as both validation and development tool

developer.intel.com/design/servers/ipmi



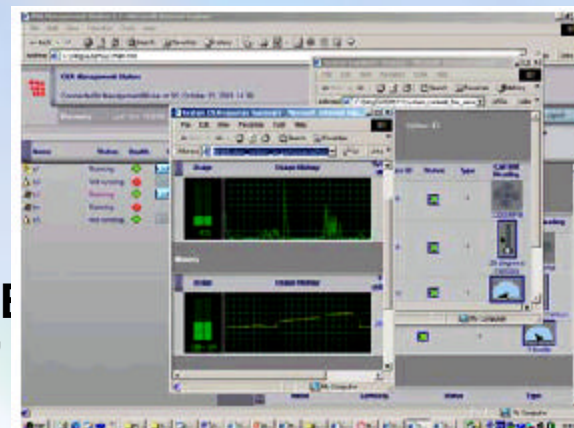
* Other names and brands may be claimed as the property of others.

Components

IPMI Software

OSA Technologies

- **Management Applications for IPMI**
 - “Remote Console” applications and IPMI Drivers
- **Firmware engineering also available**
 - SDKs for popular BMCs
 - Supports IPMI v1.5 and out-of-band access (serial, LAN)
- **IPMI LAB Thursday**
 - “Exploring Remote Manageability Building Blocks for Next Generation Servers and Devices”



**IPMI components improve TTM
and reduce design cost**




Agenda

- IPMI Update
- New system directions for IPMI
- New components for IPMI
- **IPMI future directions**
- Summary



IPMI Future Directions

Technology Transitions



Update next IDF



IPMI Future Directions

IPMI '2.0' Proposals

- Serial redirection over LAN
- Terminal mode extensions (improved 'CLI')
- ASF Alignment
 - Common authentication protocols
 - Smooths ASF to IPMI transition between desktop and sub-entry server systems
- Modular (blade) support
 - blade/chassis relationship, blade power mgmt., etc.
 - AdvancedTCA support (formerly 'CompactPCI')
- IPMI over Web
 - enabling technology for IPMI over Web (may be post 2.0)



**IPMI continues to evolve
valuable new capabilities**



Agenda

- IPMI Update
- New system directions for IPMI
- New components for IPMI
- IPMI future directions
- **Summary**



Summary

- **IPMI grows to enable competitive features across server classes**
- **Third party components improve TTM and reduce design cost**
- **IPMI continues to evolve valuable new capabilities**

developer.intel.com/design/servers/ipmi



Glossary

BMC	Baseboard Management Controller.
FRU	Field Replaceable Unit. A field replaceable component such as a board, module, fan, power supply, etc.
ICMB	Intelligent Chassis Management Bus. The ICMB provides a dedicated management bus that enables delivering IPMI messages and alerts between multiple host and peripheral chassis.
IPMB	Intelligent Platform Management Bus. Name for the architecture, protocol, and implementation of a special bus that interconnects the baseboard and chassis electronics and provides a communications media for system platform management information.
IPMI	Intelligent Platform Management Interface. IPMI defines a common, abstracted, and self-descriptive interface for platform management hardware that monitors server characteristics such as temperature, voltage, fans, power supplies, and chassis.
OOB	Out-of-Band. System platform management access that does not involve going through the OS or other software running on the main processors of the managed system.
PEF	Platform Event Filtering. A feature in IPMI that enables the BMC to generate a selectable action (e.g. power on/off, reset, send Alert, etc.) when a configurable event occurs on the management system.

Glossary

- SAF Service Availability Forum. Standards body consisting of Telco platform and software vendors that is defining RAS standards including UCMi
- SAF-HPI Server Availability Forum - Hardware Platform Interface. Name for a set of APIs and structures for representing and accessing platform management hardware.
- SAF-TE SCSI Accessed Fault-Tolerant Enclosures. SAF-TE provides a mechanism that enables RAID fault information to be sent to the hot-swap backplane via SCSI.
- SDR Sensor Data Record. SDRs provide the information that tells management software what sensors, events, management controllers, and FRU information is available from a given IPMI implementation.
- SEL System Event Log. A non-volatile storage area and associated interfaces for storing system platform event information for later retrieval.

The background of the slide features a large, detailed image of an Intel microprocessor die on the right side, with a blue and green color scheme. Swirling lines in blue, green, and yellow are visible across the background. The title text is centered and reads:

Intel Developer Forum. Spring 2003